



Laboratory for Computer Science

M.L. Dertouzos, Director

R.L. Rivest, A. Vezza, and V. Zue, Associate Directors

545 Technology Square, Cambridge, MA 02139-3539

Fax: (617) 258-8682

Phone: (617) 253-7225

internet: lynch@theory.lcs.mit.edu

January 26, 1996

Mr. Harry Koch
ESC/ENS
5 Eglin Street, Building 1704
Hanscom Airforce Base, MA 01731-2116

Dear Mr. Koch:

Please find my first quarterly R&D Status Report and Technical Summary Report for contract #C-F19628-95-C-0118 for period October 1, 1995-December 31, 1995. Copies of these reports are also being sent to our contract monitor Gary Minden and the DTIC office.

If there is any problem with the format, would you please let me know? Also, if anyone else should receive a copy of these reports, please let me know that too.

Thanks.


Sincerely,

Nancy Lynch (n)

Nancy A. Lynch
Cecil H. Green Professor
Electrical Engineering and Computer Science

cc: Gary Minden, ARPA/CSTO
DTIC Office

19960207 062



DTIC QUALITY INSPECTED 1

REPORT DOCUMENTATION PAGE		Form Approved OMB No. 0704-0188
<small>Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.</small>		
1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE 1/96	R&D Status Report 10/1/95-12/31/95
4. TITLE AND SUBTITLE Applications of the Theory of Distributed and Real Time Systems to the Development of Large-Scale Timing Based Systems		5. FUNDING NUMBERS C - F19628-95-C-0118
6. AUTHOR(S) Nancy Lynch		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Massachusetts Institute of Technology 77 Massachusetts Avenue Cambridge, MA 02138		8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Department of the Airforce Electronic Systems Center (AFMC) Hansom Air Force Base, MA 01731		10. SPONSORING / MONITORING AGENCY REPORT NUMBER
11. SUPPLEMENTARY NOTES N/A		
12a. DISTRIBUTION / AVAILABILITY STATEMENT No limits imposed on disclosure		12b. DISTRIBUTION CODE <div style="border: 1px solid black; padding: 5px; text-align: center;"> DISTRIBUTION STATEMENT A Approved for public release Distribution Unlimited </div>

This report summarizes the progress of research in the theory of distributed systems group. Members of the group have been active during this period in modeling, designing, and developing applications for concurrent systems. Our main theme is to support highly effective concurrent application designs by providing specification, quantification, and verification tools that capture behavior without limiting performance. We have developed a new "hybrid automata model" for analyzing the behaviors of systems like industrial robots or computer controlled Vehicles. We have made progress in applying our specification and automated verification methods to several complicated test problems in distributed computing. We have introduced a variety of new approaches to evaluating concurrent algorithms, among them "eventual serializability," "e-linearizability," and a novel local measure of linearizability for load balancing data structures. We are using our methods for specifying (and helping in redesigning) and a variety of new and old communication and coordination algorithms, among them Lamport's Paxos algorithm, TCP and T-TCP, and Randomized Agreement. Finally, we have shown the effectiveness of our methods in the emerging area of specifying and designing automated transit systems.

Modelling and Verification, Distributed Systems, Randomized Algorithms, Linearizability, Diffracting Trees, I/O Automata, Naming and Group Membership Services, Simulation Proofs, Hybrid Systems, Transit Systems, Vehicle Protection.

17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	15. NUMBER OF PAGES 7
			16. PRICE CODE
			20. LIMITATION OF ABSTRACT UL

R&D Status Report

Sponsored by Advanced Research Projects Agency/CSTO Applications of the theory of distributed and real time systems to the development of large scale timing-based systems.

ARPA Order No. D014

Issued by ESC/ENS under Contract #F19628-95-C-0118

Contractor Identification: Prof. Nancy Lynch, Massachusetts Institute of Technology, Laboratory for Computer Science, Cambridge, MA

Disclaimer: "The views and conclusion contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Advanced Research Projects Agency or the U. S. Government."

Progress – for period 10/1/95-12/31/95

I. Modelling and verification tools

- Lynch, Segala, Vaandrager and Weinberg have developed the "hybrid I/O automaton" model, a mathematical model based on labelled transition systems, designed for modelling and reasoning about hybrid (continuous/discrete) systems. The model includes trajectories and continuous interaction between components. The model includes composition and hiding operations, plus a notion of simulation mapping to support reasoning using levels of abstraction. Finally, it includes a notion of "receptiveness", which captures the idea that a hybrid automaton allows time to pass without bound.

This model has been presented, in preliminary form, at the 1995 DIMACS Workshop on Hybrid Systems. Lynch et al have written a preliminary paper, for inclusion in the workshop proceedings. They are currently working on a fuller version of the work, for tech report and journal publication.

For additional information, see URL <http://theory.lcs.mit.edu/tds/hybrid-model.html>.

- Pogoyants and Segala worked on adaptations of random walk theory for reasoning about randomized distributed algorithms. This is still in progress.

See URL <http://theory.lcs.mit.edu/tds/AH.html>.

- Petrov and Pogoyants carried out a machine verification, using the Larch Prover, of a complex concurrent timestamp system of Dolev and Shavit. This serves as a demonstration that our methods work well for complex examples. This is leading to the development of additional computer tools for carrying out such proofs.

See URL <http://theory.lcs.mit.edu/~petrov/CTSS.html>.

- Other projects on computer-aided verification are described in:
URL <http://theory.lcs.mit.edu/tds/cav.html> for details.
- Garland has begun the design of a Larch interface language for I/O automata. This language, together with associated tools, will facilitate reasoning about I/O automata using the Larch Prover.
See URL <http://larchj.lcs.mit.edu:8001/larch/index.html>.

II. Algorithms and impossibility results

- Lynch, Shavit, Shvartsman, and Touitou, a Ph.D. student, showed that many important classes of the highly concurrent data structures used for counting and load balancing exhibit nearly linearizable behavior for a broad range of parameters and they have completely characterized the linearizability conditions in terms of a parameter describes a local property of a low level component and that does not depend on the size of the data structure. The paper will appear in PODC96.
See URL <http://theory.lcs.mit.edu/~alex/count2.html>.
- Lynch and Rajsbaum prepared a paper explaining the mysterious Borowsky-Gafni simulation algorithm.
See URL <http://theory.lcs.mit.edu/tds/borowsky.html>.
- Shavit, Touitou and Herlihy are currently designing a highly concurrent priority queue based on the elimination tree data structures. Shavit and Della Libera are also developing a reactive version of the diffracting tree structure, that unlike existing tree implementations will exhibit an overhead that is a function of the load on the data structure.
- De Prisco and Lynch are working on highly fault tolerant distributed consensus algorithms, focusing on asynchronous systems.
See URL <http://theory.lcs.mit.edu/~robdep/research.html>.
- Shvartsman is continuing the synthesis of the latest results in the area parallel computation in the presence of failures and delays (this work was previously done jointly with Kanellakis). A monograph is in preparation.
See URL <http://theory.lcs.mit.edu/~alex/eftp.html>.

III. Applications

A. Distributed system building blocks

- Fekete, Gupta, Luchangco, Lynch and Shvartsman developed a notion of “eventually serializable object,” and demonstrated its usefulness for describing practical network name services. They also developed a general implementation, based on ideas of Ladin, Liskov and Shrira.

See URL http://theory.lcs.mit.edu/~victor_1/eventually-serializable.html.

- Following Herlihy and Wing’s well accepted notion of linearizability for specifying concurrent data structures, Lynch, Shavit, Shvartsman, and Touitou are developing a formal definition of ϵ -linearizability, a variant of linearizability that captures the notion of being “almost” linearizable, by allowing a certain fraction of concurrent operations to be out-of-order, thus facilitating data structure implementations with better scalability and load balancing properties.
- Members of the group began work on several other building blocks and their implementations, including consensus objects (implemented by Lamport’s Paxos algorithm), RAID disks, and group membership services.
- In applied work (partially done at MIT), Shvartsman summarized the emerging area of digital television system frameworks based on his consultancy to the industry.

See URL <http://theory.lcs.mit.edu/~alex/itv.html>.

- Leeb and Lynch prepared a paper modeling a steam boiler system based on a benchmark problem for a case study in formal methods for industrial applications. Preliminary results have been presented at the Dagstuhl Meeting in June 1995. This paper presents a different and improved model based on the discussions in Dagstuhl. The major improvement is that the new approach guarantees safety in the presence of faults.

See URL <http://theory.lcs.mit.edu/tds/boiler.html>.

B. Transit

- Weinberg and Lynch have a nearly-complete version of an analysis, using invariants and simulations, of a simple vehicle deceleration maneuver. This will appear in H.B. Weinberg’s M.S. thesis.

See URL <http://theory.lcs.mit.edu/~hbw/decel.html>.

- Lynch, Weinberg and Delisle presented some preliminary ideas on modelling and analyzing separate vehicle protection (VP) subsystems, in the DIMACS-95 Workshop. They are currently writing a paper for the workshop proceedings.

See URL <http://theory.lcs.mit.edu/~hbw/prot.html>.

- Lynch has written a preliminary report on the analysis of a simple acceleration maneuver, using levels of abstraction.

See URL <http://theory.lcs.mit.edu/tds/three-level.html>.

- Lynch presented a survey of our group's work so far on modelling and analyzing transit systems, also at the DIMACS-95 Workshop. She has written a paper for the workshop proceedings.

See URL <http://theory.lcs.mit.edu/tds/transit-survey.html>.

C. Communication

- Smith produced a good outline of a correctness proof for TCP, and began a correctness proof for an optimization, T-TCP.

See URL <http://theory.lcs.mit.edu/~mass/comm.html>.

D. Probabilistic Systems

- Pogosyants and Segala have produced preliminary proofs for many of the properties of the (randomized) Aspnes-Herlihy consensus protocol.

See URL <http://theory.lcs.mit.edu/tds/AH.html>.

These proofs are based on Segala's general semantic model for probabilistic systems. See

URL <http://theory.lcs.mit.edu/tds/probability.html>.

Changes in Key Personnel

- Three new graduate students joined in September.
- Anya Pogosyants, one of the PhD students working on this project, was killed on December 15 in a car crash on an icy road in Vermont.
- Alex Shvartsman's major collaborator, Prof. Paris Kanellakis of Brown University, was killed in the Colombian plane crash, later in December.
- Dr. Steve Garland has joined the project as a collaborator to help construct computer tools for simulating and verifying distributed algorithms.

Travel and Relevant Meetings

- Lynch and Weinberg met with system developers at Raytheon to discuss the modelling/verification of a Raytheon vehicle protection system.
- Lynch and Weinberg attended the DIMACS workshop on hybrid systems in October. Three papers from our group were presented.
See URL's <http://theory.lcs.mit.edu/tds/transit-survey.html>,
<http://theory.lcs.mit.edu/tds/LSVW.html>, and
<http://theory.lcs.mit.edu/~hbw/protftp.html>.
They learned about the control theorists' view of the hybrid systems area.
- Several members of Prof. Shankar Sastry's PATH group at Berkeley visited us in October for discussions of the PATH system and its modelling/verification.
- Lynch travelled to Bellcore in December to meet with designers of a network administration system. They discussed definition of building blocks (in particular, group membership services, replicated data services and communication services) needed for the Bellcore system.
- Alex Shvartsman participated in the SPIE Conference on Integration Issues in Large Scale Media Delivery Systems, where he gave a talk on integration frameworks for digital video networks.

Budget Information

Current funding is sufficient for the current fiscal year. The next fiscal year's funding requirement at current, anticipated levels is \$190,500.

10/1/95-12/31/95 Publications

- [1] Y. Afek, B. Awerbuch, E. Gafni, Y. Mansour, A. Rosen, and N. Shavit. Slide: the key to polynomial end-to-end communication. To appear in *Journal of Algorithms*.
- [2] J.F. Buss, P.C. Kanellakis, P. L. Ragde, and A. Shvartsman. Parallel algorithms with processor failures and delays. *Journal of Algorithms*, 1996. To appear.
- [3] Alan Fekete, David Gupta, Victor Luchangco, Nancy Lynch, and Alex Shvartsman. Eventually-serializable data services. In *Proceedings of the Fifteenth Annual ACM Symposium on Principles of Distributed Computing*, Philadelphia, PA, May 1996. To appear.

- [4] M. P. Herlihy and N. Shavit. The topological structure of asynchronous computability. To appear as a Brown University TR. Also, submitted for publication.
- [5] Maurice Herlihy, Nir Shavit, and Orli Waarts. Low contention linearizable counting. To appear in *Distributed Computing*.
- [6] P.C. Kanellakis, D. Michailidis, and A. Shvartsman. Controlling memory access in efficient fault-tolerant parallel algorithms. *Nordic Journal Computer Science*, 2:146–180, 1995.
- [7] Gunter Leeb and Nancy Lynch. Proving safety properties of the steam boiler controller: Formal methods for industrial applications, a case study, January 1996. Submitted for publication. Presented at the *Methods for Semantics and Specification*, International Conference and Research Center for Computer Science, Schloss, Dagstuhl, Germany, June 1995, as “Using Timed Automata for the Steam Boiler Controller Problem.”
- [8] Nancy Lynch. A three-level analysis of a simple acceleration maneuver, with uncertainties. Manuscript. WWW URL=<http://theory.lcs.mit.edu/three-level.html>.
- [9] Nancy Lynch. Modelling and verification of automated transit systems, using timed automata, invariants and simulations. In *DIMACS Workshop on Verification and Control of Hybrid Systems*, October 1995. To appear in R. Alur, T. Henzinger, and E. Sontag, editors, *Hybrid Systems III*, Lecture Notes in Computer Science, Springer-Verlag. Also, to appear as MIT/LCS/TM-545.
- [10] Nancy Lynch and Sergio Rajsbaum. On the Borowsky-Gafni simulation algorithm. In *Proceedings of the Fifteenth Annual ACM Symposium on Principles of Distributed Computing*, Philadelphia, PA, May 1996. Short version. To appear.
- [11] Nancy Lynch and Roberto Segala. A comparison of simulation techniques and algebraic techniques for verifying concurrent systems. *Formal Aspects of Computing*, 7(3):231–265, 1995.
- [12] Nancy Lynch, Roberto Segala, Frits Vaandrager, and H.B. Weinberg. Hybrid I/O automata. In *DIMACS Workshop on Verification and Control of Hybrid Systems*, October 1995. To appear in R. Alur, T. Henzinger, and E. Sontag, editors, *Hybrid Systems III*, Lecture Notes in Computer Science, Springer-Verlag. Also, to appear as MIT/LCS/TM-544.
- [13] Nancy Lynch, Nir Shavit, Alex Shvartsman, and Dan Touitou. Counting networks are practically linearizable. In *Proceedings of the Fifteenth Annual ACM Symposium on Principles of Distributed Computing*, Philadelphia, PA, May 1996. To appear.

- [14] N. Shavit and D. Touitou. Software transactional memory. Submitted for publication. Also, to be Massachusetts Institute of Technology, Laboratory for Computer Science, MIT/LCS/TR-675.
- [15] N. Shavit, E. Upfal, and A. Zemach. A steady state analysis of diffracting trees. Manuscript. Submitted for conference publication.
- [16] Alex Shvartsman. Integrating distributed multimedia systems and interactive television networks. In *SPIE Conf. on Integration of Large Scale Media Delivery Systems*, volume 2615, Philadelphia, October 1995.
- [17] H.B. Weinberg, Nancy Lynch, and Norman Delisle. Verification of automated vehicle protection systems. In *DIMACS Workshop on Verification and Control of Hybrid Systems*, October 1995. To appear in R. Alur, T. Henzinger, and E. Sontag, editors, *Hybrid Systems III*, Lecture Notes in Computer Science, Springer-Verlag.